

# **“KernelSizeTuningGuide”**

## **Numerical examination result (and tips)**

Oct 27th, 2006

IKEDA, Munehiro

m-ikeda@ds.jp.nec.com



# Contents

- **Background**
- **Examination scope**
- **Effects of each config-item**
  - Targets and conditions
  - Examination result
- **Kernel overall**
  - Targets and conditions
  - Examination result
- **Further works**
  - Appendix 1 : Target config-items and effects
  - Appendix 2 : default/small config-set
  - Appendix 3 : Tips



# Background

- Our activity :  
KconfigSize tool  
<http://tree.celinuxforum.org/CelfPubWiki/KernelConfigWeight>
    - A tool to automatically examine “impact” by config-items against kernel static size and RAM usage.
  - CELF Public Wiki :  
“Linux Kernel Size Configuration Guide”  
<http://tree.celinuxforum.org/CelfPubWiki/KernelSizeTuningGuide>
    - The page describes worth informations to reduce kernel static size and RAM usage. Published by Mr.Bird.
    - Table in “4.3 Kernel Configuration Options” shows configuration options which are effective to reduce kernel size.  
(Config-items introduced by Linux-Tiny are included.)
      - Linux-Tiny ... a patch-set to reduce kernel size.
    - The table proposes “default” and “small” settings.
      - small ... recommended setting to reduce size.
- ➔ KconfigSize tool can numerically examine the difference between “default” and “small” .



# Examination scope

- To contribute numerical data to  
“4.3 Kernel Configuration Options”  
of “Linux Kernel Size Configuration Guide”
    - Examining effects of each config-item  
for static size and RAM usage
      - How kernel static size is smaller when set a config-item  
“small” than “default” .
      - How RAM usage is less in the same way above.
    - Examining static size and RAM usage of overall kernel,  
by measurement 4 kernels below.
      - Kernel sources : vanilla / Linux-Tiny applied
      - Config-set : “default” / “small”
- \* Definition of “RAM usage” is amount of RAM which BUDDY recognizes as in-use just after boot of kernel (= just before execution of userland).  
It should be “just-for-reference” because there is reproducibility error from our experiences.



# Effects of each config-item : targets and conditions

- **Target config-items**

- Config-items listed in “4.3 Kernel Configuration Options” of “Linux Kernel Size Configuration Guide” , and applied little modifications as below.
  - Used new symbol names if whose names had been changed.
  - Omitted items which had been vanished.
  - Omitted items which occurred compilation failure when enabled.
  - Added several items which may have significant effect.

➔Examined 41 items

\* (All items are listed in “Appendix 1” )

- **Environment**

- Kernel 2.6.16.19 + Linux-Tiny patches
- arch : i386
- Platform for RAM usage measurement : qemu 0.6.2, mem=128MB (on Pentium IV 3.4GHz, mem=1.0GB)
- gcc : 3.3.5
- binutils : 2.15



# Effects of each config-item : examination result (1)

- **Effects for static size**

- 15 items in the table below have effect > 10KB for bzImage size.

- \* Results of all items are listed in “Appendix 1” .

| CONFIG option               | setting |    | vanilla<br>/tiny | difference[byte] |          | note                 |
|-----------------------------|---------|----|------------------|------------------|----------|----------------------|
|                             | from    | to |                  | vmlinux          | bzImage  |                      |
| CONFIG_PCI                  | Y       | N  | vanilla          | -1,217,753       | -463,200 |                      |
| CONFIG_CC_OPTIMIZE_FOR_SIZE | N       | Y  | vanilla          | -487,424         | -167,904 |                      |
| CONFIG_PRINTK               | Y       | N  | vanilla          | -365,530         | -142,176 | introduced from tiny |
| CONFIG_KALLSYMS             | Y       | N  | vanilla          | -235,189         | -130,592 |                      |
| CONFIG_SCSI                 | Y       | N  | vanilla          | -284,965         | -114,720 |                      |
| CONFIG_IDE                  | Y       | N  | vanilla          | -175,576         | -69,568  |                      |
| CONFIG_SMP                  | Y       | N  | vanilla          | -177,845         | -64,064  |                      |
| CONFIG_MODULES              | Y       | N  | vanilla          | -369,445         | -44,608  |                      |
| CONFIG_NFS_FS               | Y       | N  | vanilla          | -63,186          | -28,384  |                      |
| CONFIG_SYSCTL               | Y       | N  | vanilla          | -84,593          | -28,000  |                      |
| CONFIG_NFSD                 | Y       | N  | vanilla          | -56,816          | -24,256  |                      |
| CONFIG_BUG                  | Y       | N  | vanilla          | -43,287          | -14,848  | introduced from tiny |
| CONFIG_SYSVIPC              | Y       | N  | vanilla          | -27,365          | -12,640  |                      |
| CONFIG_SWAP                 | Y       | N  | vanilla          | -19,266          | -11,136  |                      |
| CONFIG_XATTR                | Y       | N  | tiny             | -21,517          | -10,816  | see “Tips”           |



# Effects of each config-item : examination result (2)

- **Effects for RAM usage (for reference)**
  - 12 items in the table below have effect > 100KB for RAM usage.
    - \* Results of all items are listed in “Appendix 1” .

| CONFIG option           | setting |    | vanilla<br>/tiny | difference[byte]<br>used mem | note |
|-------------------------|---------|----|------------------|------------------------------|------|
|                         | from    | to |                  |                              |      |
| CONFIG_NFSD             | Y       | N  | vanilla          | -454,656                     |      |
| CONFIG_SCSI             | Y       | N  | vanilla          | -315,392                     |      |
| CONFIG_PCI              | Y       | N  | vanilla          | -249,856                     |      |
| CONFIG_NR_LDISCS        | 16      | 2  | tiny             | -163,840                     |      |
| CONFIG_UID16            | Y       | N  | vanilla          | -163,840                     |      |
| CONFIG_BLK_DEV_LOOP     | Y       | N  | vanilla          | -163,840                     |      |
| CONFIG_KMOD             | Y       | N  | vanilla          | -159,744                     |      |
| CONFIG_MODULES          | Y       | N  | vanilla          | -147,456                     |      |
| CONFIG_IOSCHED_DEADLINE | Y       | N  | vanilla          | -147,456                     |      |
| CONFIG_MAX_USER_RT_PRIO | 100     | 5  | tiny             | -143,360                     |      |
| CONFIG_IP_PNP_RARP      | Y       | N  | vanilla          | -135,168                     |      |
| CONFIG_AUDIT_BOOTMEM    | Y       | N  | tiny             | -135,168                     |      |

- \* CONFIG\_MEMPOOL(mempool-shrink.patch) showed significant effect for RAM usage in our experiment before, though it is out of scope of this examination.



# Kernel overall : targets and conditions

- **Target kernels :**  
**4 kernels from 2 kinds of sources and 2 kinds of config-sets.**
  - Kernel sources
    - vanilla: source code from kernel.org
    - tiny : source code which is applied Linux-Tiny on vanilla
  - config-set
    - “default”
      - Based on defconfig for i386, and applied modifications as below.
        - » Changed all “=m” into “=y”.  
... to measure sizes of vmlinux/bzImage(s)
        - » Enabled CONFIG\_SERIAL\_8250\_CONSOLE=y  
... because RAM usage measurement platform uses serial console.
    - “small”
      - Based on “default”, and config-items which were target items of “effects of each config-item” were set one which made static size smaller.
- \* Please see “Appendix 2” for detail of config-set.
- **Environment : same as “effects of each config-item” (see p5)**







# Kernel overall : examination result (1)

- **Static size**

- vanilla (bzImage)
  - default : 2.0MB  $\Rightarrow$  small : 900KB
- tiny (bzImage)
  - default : 2.0MB  $\Rightarrow$  small : 860KB




## Kernel static size[BYTE]

| kernel  | config-set | vmlinux   | bzImage   | \$(size vlinux) |         |         |
|---------|------------|-----------|-----------|-----------------|---------|---------|
|         |            |           |           | text            | data    | bss     |
| vanilla | default    | 5,267,762 | 2,019,678 | 3,606,842       | 561,724 | 186,972 |
|         | small      | 2,190,346 | 903,209   | 1,547,095       | 159,976 | 64,276  |
| tiny    | default    | 5,267,871 | 2,014,684 | 3,603,061       | 562,236 | 185,788 |
|         | small      | 2,112,179 | 863,484   | 1,478,058       | 159,836 | 57,844  |



# Kernel overall : examination result (2)

- **RAM usage (for reference)**
  - vanilla (used mem)
    - default : 3.8MB  $\Rightarrow$  small : 2.9MB
  - tiny (used mem)
    - default : 4.0MB  $\Rightarrow$  small : 3.0MB



## RAM usage [BYTE]

| kenrel  | config-set | used      | free        | total       |
|---------|------------|-----------|-------------|-------------|
| vanilla | default    | 3,829,760 | 124,260,352 | 128,090,112 |
|         | small(+)   | 2,920,448 | 127,713,280 | 130,633,728 |
| tiny    | default    | 3,977,216 | 124,112,896 | 128,090,112 |
|         | small(+)   | 3,022,848 | 127,692,800 | 130,715,648 |

\* small(+) ... “small” + CONFIG\_PRINTK=y, CONFIG\_IDE=y  
necessary for booting and measurement on our environment.



# Further works

- To upload these data to CELF Public WiKi.
- To expand data
  - Other kernel versions
  - Archs other than i386
  - Config-items which are out of scope of this examination
- ➔ Automatic examination on CELF Test Lab. is preferable.
- The examination tool is still needed to be improved little bit...



## Appendix 1 : Target config-items and effects (1/3)

| CONFIG option              | setting |    | difference[byte] |         |          | note   |
|----------------------------|---------|----|------------------|---------|----------|--|
|                            | from    | to | vmlinux          | bzImage | used mem |  |
| CONFIG_BASE_FULL           | Y       | N  | -4096            | -32     | -16384   | changed from CONFIG_CORE_SMALL (which was for 2.6.10 in core-small.patch)                              |
| CONFIG_NET_SMALL           | N       | Y  | 0                | 32      | 16384    |  |
| CONFIG_KMALLOC_ACCOUNTING  | N       | N  | -                | -       | -        | compilation failure when enabled (mm/slab.c : __do_kmalloc())<br>kmem_cache has no member of "cs_size" |
| CONFIG_AUDIT_BOOTMEM       | Y       | N  | 0                | -128    | -135168  |  |
| CONFIG_MEASURE_INLINES     | Y       | N  | 0                | 0       | 0        | changed from CONFIG_DEPRECATED_INLINES (typo?)   |
| CONFIG_PRINTK              | Y       | N  | -365530          | -142176 | -        | mem measurement uses printk  |
| CONFIG_BUG                 | Y       | N  | -43287           | -14848  | 12288    |  |
| CONFIG_ELF_CORE            | Y       | N  | -4424            | -2880   | -8192    |  |
| CONFIG_PROC_KCORE          | Y       | N  | -4467            | -3328   | 20480    |  |
| CONFIG_AIO                 | Y       | N  | -9329            | -7328   | -16384   |  |
| CONFIG_XATTR               | Y       | N  | -21517           | -10816  | 77824    | see "Tips"   |
| CONFIG_FILE_LOCKING        | Y       | Y  | -                | -       | -        | compilation failure (see "Tips")   |
| CONFIG_DIRECTIO            | -       | -  | -                | -       | -        | obsolete (was 2.6.10 in direct-io-core.patch)  |
| CONFIG_MAX_SWAPFILES_SHIFT | 5       | 0  | 0                | -96     | -4096    |  |



## Appendix 1 : Target config-items and effects (2/3)

(cont'd)

| CONFIG option                      | setting |    | difference[byte] |         |          | note                                 |
|------------------------------------|---------|----|------------------|---------|----------|--------------------------------------|
|                                    | from    | to | vmlinux          | bzImage | used mem |                                      |
| CONFIG_KALLSYMS                    | Y       | N  | -235189          | -130592 | -4096    |                                      |
| CONFIG_SHMEM                       | Y       | N  | -14233           | -7360   | -12288   |                                      |
| CONFIG_SWAP                        | Y       | N  | -19266           | -11136  | 8192     |                                      |
| CONFIG_SYSVIPC                     | Y       | N  | -27365           | -12640  | -4096    | changed from CONFIG_SYSV_IPC (typo?) |
| CONFIG_POSIX_MQUEUE                | Y       | N  | -9681            | -3936   | -4096    |                                      |
| CONFIG_SYSCTL                      | Y       | N  | -84593           | -28000  | -57344   |                                      |
| CONFIG_LOG_BUF_SHIFT               | 15      | 12 | 0                | 32      | 20480    | see "Tips"                           |
| CONFIG_UID16                       | Y       | N  | -4408            | -3040   | -163840  |                                      |
| CONFIG_CC_OPTIMIZE_FOR_SIZE        | N       | Y  | -487424          | -167904 | 12288    |                                      |
| CONFIG_MODULES                     | Y       | N  | -369445          | -44608  | -147456  |                                      |
| CONFIG_KMOD                        | Y       | N  | -4959            | -4128   | -159744  |                                      |
| CONFIG_PCI                         | Y       | N  | -1217753         | -463200 | -249856  |                                      |
| <del>CONFIG_XIP_KERNEL</del>       | -       | -  | -                | -       | -        | not supported for i386 (arch depend) |
| <del>CONFIG_MAX_RESERVE_AREA</del> | -       | -  | -                | -       | -        | ?? (arch depend?)                    |
| CONFIG_BLK_DEV_LOOP                | Y       | N  | -14292           | -7488   | -163840  |                                      |
| CONFIG_BLK_DEV_RAM                 | Y       | N  | -10237           | -4512   | -40960   |                                      |
| CONFIG_BLK_DEV_RAM_COUNT           | 16      | 2  | 0                | -32     | -28672   |                                      |



## Appendix 1 : Target config-items and effects (3/3)

(cont'd)

| CONFIG option           | setting |    | difference[byte] |         |          | note  |
|-------------------------|---------|----|------------------|---------|----------|---|
|                         | from    | to | vmlinux          | bzImage | used mem |   |
| CONFIG_IOSCHED_AS       | Y       | N  | -10982           | -7680   | -8192    |   |
| CONFIG_IOSCHED_DEADLINE | Y       | N  | -6107            | -3008   | -147456  |   |
| CONFIG_IOSCHED_CFQ      | Y       | N  | -12393           | -6880   | -28672   |   |
| CONFIG_IP_PNP           | Y       | N  | -14156           | -7168   | -16384   |   |
| CONFIG_IP_PNP_DHCP      | Y       | N  | -101             | -224    | -4096    |   |
| CONFIG_IP_PNP_BOOTP     | Y       | N  | 0                | 0       | -20480   | Addedby IKEDA   |
| CONFIG_IP_PNP_RARP      | Y       | N  | -4190            | -2784   | -135168  | Added by IKEDA  |
| CONFIG_IDE              | Y       | N  | -175576          | -69568  | -        | mem measurement environment uses IDE as a root device |
| CONFIG_SCSI             | Y       | N  | -284965          | -114720 | -315392  |   |



## Appendix 2 : default/small config-set (1/3)

| CONFIG option              | vanilla |       | tiny    |       | note                             |
|----------------------------|---------|-------|---------|-------|----------------------------------|
|                            | default | small | default | small |                                  |
| CONFIG_BASE_FULL           | Y       | N     | Y       | N     |                                  |
| CONFIG_NET_SMALL           | –       | –     | N       | Y     |                                  |
| CONFIG_KMALLOC_ACCOUNTING  | –       | –     | N       | N     | compilation failure when enabled |
| CONFIG_AUDIT_BOOTMEM       | –       | –     | N       | N     | worth for development            |
| CONFIG_MEASURE_INLINES     | –       | –     | N       | N     | worth for development            |
| CONFIG_PRINTK              | Y       | N     | Y       | N     |                                  |
| CONFIG_BUG                 | Y       | N     | Y       | N     |                                  |
| CONFIG_ELF_CORE            | Y       | N     | Y       | N     |                                  |
| CONFIG_PROC_KCORE          | Y       | N     | Y       | N     |                                  |
| CONFIG_AIO                 | –       | –     | Y       | N     |                                  |
| CONFIG_XATTR               | –       | –     | Y       | N     | see “Tips”                       |
| CONFIG_FILE_LOCKING        | –       | –     | Y       | N     | see “Tips”                       |
| CONFIG_DIRECTIO            | –       | –     | –       | –     |                                  |
| CONFIG_MAX_SWAPFILES_SHIFT | –       | –     | 5       | 0     |                                  |
| CONFIG_NR_LDISCS           | –       | –     | 16      | 2     |                                  |



(cont'd)

| CONFIG option                      | vanilla |       | tiny    |       | note   |
|------------------------------------|---------|-------|---------|-------|--|
|                                    | default | small | default | small |  |
| CONFIG_KALLSYMS                    | Y       | N     | Y       | N     |  |
| CONFIG_SHMEM                       | Y       | N     | Y       | N     |  |
| CONFIG_SWAP                        | Y       | N     | Y       | N     |  |
| CONFIG_SYSVIPC                     | Y       | N     | Y       | N     |  |
| CONFIG_POSIX_MQUEUE                | Y       | N     | Y       | N     |  |
| CONFIG_SYSCTL                      | Y       | N     | Y       | N     |  |
| CONFIG_LOG_BUF_SHIFT               | 15      | 12    | 15      | 12    | see “Tips”   |
| CONFIG_UID16                       | Y       | N     | Y       | N     |  |
| CONFIG_CC_OPTIMIZE_FOR_SIZE        | N       | Y     | N       | Y     |  |
| CONFIG_MODULES                     | Y       | N     | Y       | N     |  |
| CONFIG_KMOD                        | Y       | (N)   | Y       | (N)   | depends on MODULES   |
| CONFIG_PCI                         | Y       | N     | Y       | N     |  |
| <del>CONFIG_XIP_KERNEL</del>       | –       | –     | –       | –     |  |
| <del>CONFIG_MAX_RESERVE_AREA</del> | –       | –     | –       | –     |  |
| CONFIG_BLK_DEV_LOOP                | N       | N     | N       | N     |  |
| CONFIG_BLK_DEV_RAM                 | N       | N     | N       | N     |  |
| CONFIG_BLK_DEV_RAM_COUNT           | (16)    | (2)   | (16)    | (2)   | default 16 if !BLK_DEV_RAM<br>“2” can be set only when<br>CNOFIG_BLK_DEV_RAM=y |





(cont'd)

| CONFIG option           | vanilla |       | tiny    |       | note                   |
|-------------------------|---------|-------|---------|-------|------------------------|
|                         | default | small | default | small |                        |
| CONFIG_IOSCHED_AS       | Y       | N     | Y       | N     |                        |
| CONFIG_IOSCHED_DEADLINE | Y       | N     | Y       | N     |                        |
| CONFIG_IOSCHED_CFQ      | Y       | N     | Y       | N     |                        |
| CONFIG_IP_PNP           | N       | N     | N       | N     |                        |
| CONFIG_IP_PNP_DHCP      | N       | (N)   | N       | (N)   | depends on IP_PNP_DHCP |
| CONFIG_IP_PNP_BOOTP     | N       | (N)   | N       | (N)   | depends on IP_PNP_DHCP |
| CONFIG_IP_PNP_RARP      | N       | (N)   | N       | (N)   | depends on IP_PNP_DHCP |
| CONFIG_IDE              | Y       | N     | Y       | N     |                        |
| CONFIG_SCSI             | Y       | N     | Y       | N     |                        |
| CONFIG_SMP              | Y       | N     | Y       | N     |                        |



Appearing below are “tips” if that helps.

- **CONFIG\_XATTR**

- When XATTR=n, EXT3\_FS\_XATTR=y causes compilation failure.  
( EXT3\_FS\_XATTR=y in defconfig )  
=> Need to add “EXT3\_FS\_XATTR depends on XATTR” on fs/Kconfig

- **CONFIG\_FILE\_LOCKING**

- When FILE\_LOCKING=n, NFSD=y causes compilation failure.  
( NFSD=y in defconfig )  
=> Need to add “NFSD depends on FILE\_LOCKING” on fs/Kconfig



- **CONFIG\_LOG\_BUF\_SHIFT**

- LOG\_BUF\_SHIFT can't be set unless DEBUG\_KERNEL=y because “if DEBUG\_KERNEL” in lib/Kconfig.  
Please make sure to disable below items to make kernel compact, which are “depends on DEBUG\_KERNEL” and “default y” .
  - DETECT\_SOFTLOCKUP
  - DEBUG\_PREEMPT
  - DEBUG\_MUTEXES
  - FORCED\_INLINING

- **CONFIG\_PCI**

- Disable ACPI in advance if you want to disable PCI because ACPI selects PCI.

- **CONFIG\_PRINTK**

- syslogd needs PRINTK=y. You shouldn't launch syslogd on kernel configured as PRINTK=n.



Thank you !